ALCF Early Science Program

ESP Kick-Off Workshop
Project Plan Presentation

Materials Design and Discovery: Catalysis and Energy Storage

PI: Larry Curtiss

Presenter: Nichols A. Romero

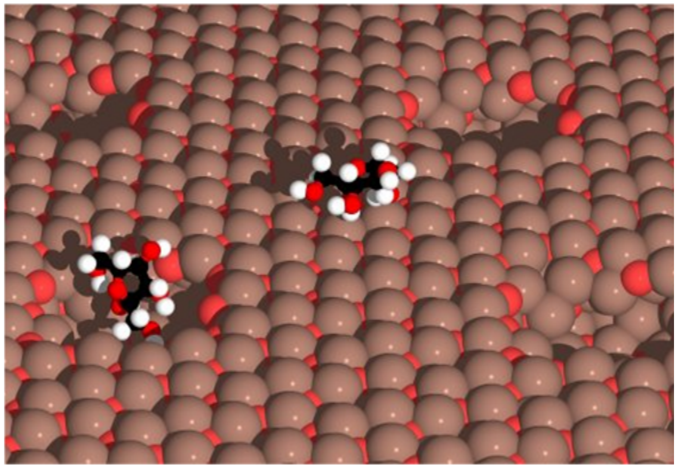

October 18-19, 2010

Project Overview

Material science problems critical to DOE's energy security mission.

- **Electronic structure calculations for studying:**
  - Biomass conversion
  - Electric energy interfaces (Li-ion batteries)
  - Lithium-air batteries
  - Catalysis of transition metal (Ag, Au, Pt) nanoparticles
- **Cross-cuts multiple efforts at ANL:**
  - Energy Frontier Research Centers (EFRCs)
    - Institute for Atom-efficient Chemical Transformations (IACT)
    - Center for Electrical Energy Storage (CEES)
  - LDRD Director's Grand Challenge
  - Center for Nanoscale Materials (CNM)

Codes: GPAW

Scientific Field: Chemistry and Material Science

Project Overview (cont'd)

- **Material science problems are hard, really hard:**
  - Li-ion battery specific energy – 900 W-h/kg
  - Octane specific energy – 13,000 W-h/kg
  - Need to optimize: energy density, power density, affordability, temperature resilient
- **Materials discovery requires solving rational design problem:**

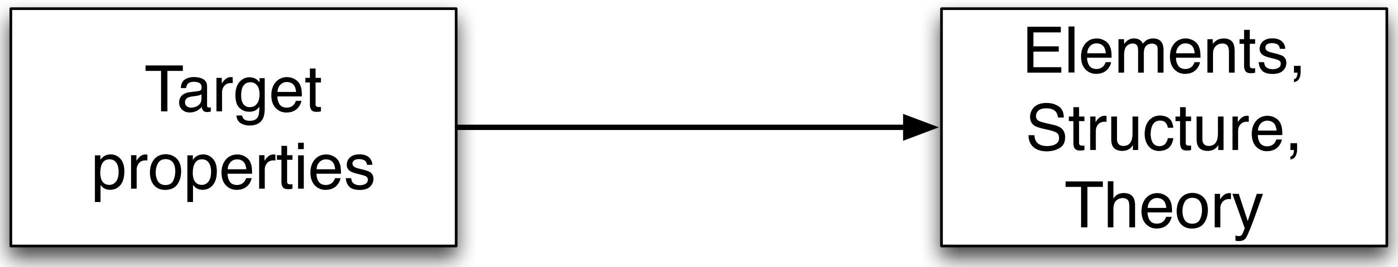

- **Some things to note:**
  - Higher resolution (e.g. denser grids) with current methods on existing problems is useless.
  - Chemically accurate methods are very expensive (N^4, N^5, N^7, exp(N)).
  - Cheap methods (N^X, where $1 < x < 3$) often fail for the most important materials (strong-correlated).
  - Search through exponential large configure space of materials is non-trivial.

Project Overview (cont'd)

- **Mira will allow calculations on:**
  - Larger systems than on Blue Gene/P
  - High-throughput computing on top of large MPI problems
  - Loosely-coupled images for computing energy barrier (NEB)
  - Some higher accuracy electronic structure calculations with the more advanced methods (e.g. QMC). [not part of this ESP, but possible future work]

Computational Approach, Numerical Methods

- **Kohn-Sham (KS) Density Functional Theory (DFT)**
  - mean-field theory using single-particle wave functions
  - method formally scales as $O(N^3)$; double the problem size requires eight times FLOPs
  - number of algorithms in the scientific community is dizzying
    - representations for wavefunctions
    - boundary conditions
    - methods for metals or insulators (or both)
    - many ways to minimize the Kohn-Sham energy functional
- **GPAW is an implementation of the PAW method**
  - pseudowavefunctions on a uniform real-space grid
  - Projector Augmented Wave method (PAW) accurately treats core electrons
  - Written in Python and C. Less than 10% of the time is spent in Python
- **KS equation solved through self-consistent field (SCF) iterations**
  - Non-linear sparse eigenvalue problem for valence electrons
  - Iterative diagonalization using Residual Minimization Method - Direct Inversion in the Iterative Subspace (RMM-DIIS), a Lanczos type method with no deflation
  - Pulay mixing of the charge density

Computational Approach, Numerical Methods (cont'd)

- **RMM-DIIS method operations. N_g, N_p, N_b are total number of grid points, projectors (~atoms), and bands (~electrons), respectively. $N_g >> N_p > N_b$.**
  - Solve the Poisson equation [$O(N_g)$]
  - Exchange-correlation energy and potential [$O(N_g)$]
  - Density mixing [$O(N_g)$]
  - $H\Psi$ are *sparse operations* [$O(N_b N_g)+O(N_b N_p)$]
  - Constructing the pseudo-density [$O(N_b N_g)$]
  - Matrix elements and orthogonalization are *dense operations* (DGEMM). [$O(N_b^2 N_g)$]
  - Cholesky decomposition (ScaLAPACK) [$O(N_b^3)$]
  - Subspace diagonalization (ScaLAPACK) [$O(N_b^3)$]
  - This leads to unusual Amdahl law bottlenecks.

Parallelism and Existing Implementation

- **Five simultaneous layers of parallelization are possible**
  - k-points, spins, bands, domains, and images
  - only two-three layers of parallelization used in practice (bands, domains, spins/images)
  - projectors are non-trivial to load-balance, would require global task scheduling (GTS)
- **MPI everywhere (no threads)**
- **HDF5 for I/O**
- **Current Performance/Scalability**
  - GPAW single-point energy LDA/PBE calculation scales to 8-racks of Intrepid
    - 65(85) % strong scaling efficiency (latter if you omit one-time overhead costs)
    - 20% peak of single-core performance
    - 1 minute per SCF step
    - $O(N_b^3)$ operation handled by ScaLAPACK are difficult to scale.
  - Loading Python modules and dynamic libraries pose I/O bottleneck at 8-racks
    - Present work around is ramdisk
    - William Scullin (ALCF) working on parallel loader built on top of SPI. Broadcast python libraries and dynamic libraries over tree network. Work in progress.
    - LLNL have reported similar problems on their Blue Gene/P.
    - Problem reproduced by N. Smeds (IBM).

Library and Tool Dependencies

- **Libraries**
  - NumPy: numerical Python for manipulating arrays
  - Atomic simulation environment (ASE) Python for setting up calculations
  - ESSL for optimized BLAS
    - heavy DGEMM use with alpha.neq.1 and beta.eq.1
  - LAPACK, ScaLAPACK for dense linear algebra and parallel dense linear algebra, respectively.
  - HDF5 needed for restarting calculation with wavefunctions
- **Tools**
  - HPM library for accessing performance counters
    - Current implementation requires timer to be called on MPI_COMM_WORLD.
  - Tuning and Analysis Utilities (TAU) for performance analysis of Python, C, and MPI.
  - ParaView (low priority and non-essential)
    - need help with visualization charge density, thanks to the MCS sci-vis team

Anticipated Modifications for Blue Gene/Q

- **Implement fine-grain parallelism**
  - threaded ESSL for large DGEMM
  - OpenMP for C kernels
- **Threaded parallel dense linear algebra library**
  - ScaLAPACK not ready for petascale, looking into Elemental as a potential alternative.
  - Jack Poulson (Elemental developer) identified performance issues in Blue Gene/P MPI software stack; hopefully this will be resolved for Blue Gene/Q.
    - MPI_Allgather, MPI_Reducescatter
    - Ability to force MPI to use a particular algorithm needed. Discussions with B. Smith (IBM).
  - Elemental is not complete and may not be complete for another year. ScaLAPACK would have to do for now.
  - Basic research in dense diagonalization algorithms is needed.
    - Heavy-use of DGEMV. Lots of load imbalance.
    - Previous guess for eigenvectors/eigenvalues not used to accelerate convergence
- **Performance and scaling needed to run proposed problems on Mira:**
  - Nudge elastic band (NEB) calculations are embarrassingly parallel.
  - Twice larger problems require scaling out eight times as far, still at 1 min per SCF. Big trouble if ScaLAPACK behaves poorly.

Plan for Next 6 Months Effort

- **Post-doc will be interviewed in early November**
- **Jussi Enkovaara (CSC, Ltd.) has developed OpenMP branch of GPAW**
  - Careful performance analysis of computer kernels needed.
  - MPI everywhere still has better performance than OpenMP.
- **Detailed performance measurements on Blue Gene/P**
  - Collecting data using TAU for the last couple of years
  - Aware of bottlenecks
- **Introduce OpenMP in main code kernels:**
  - Localized functions – integrate and add; critical to maximize FLOP rate on BG/P
  - Interpolator and restrictor for wavefunctions and charge density
  - Poisson solver
- **Some Python code will require migration to C**
- **A. R. Mamidala (IBM) taking a closer look at ScaLAPACK performance**
- **W. Scullin (ALCF) will continue to work on parallel loader**
- **Use projections to estimate performance on BG/Q**